




Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.



The economic impact of lockdowns: A theoretical assessment[☆]

Gabriele Camera^{a,b,*}, Alessandro Giorfré^c

^a Economic Science Institute, Chapman University, One University dr., Orange, CA 92866, United States of America

^b DSE, University of Bologna, Italy

^c DISEI, University of Florence, Italy

ARTICLE INFO

Article history:

Received 8 February 2021

Received in revised form 1 June 2021

Accepted 8 July 2021

Available online 18 July 2021

Manuscript handled by Editor Andres Carvajal

Keywords:

Decentralized markets

Random matching

Contagion

Nonpharmaceutical interventions

ABSTRACT

The sudden appearance of the SARS-CoV-2 virus and the onset of the COVID-19 pandemic triggered extreme and open-ended “lockdowns” to manage the disease. Should these drastic interventions be the blueprint for future epidemics? We construct an analytical framework, based on the theory of random matching, which makes explicit how epidemics spread through economic activity. Imposing lockdowns by assumption not only prevents contagion and reduces healthcare costs, but also disrupts income-generation processes. We characterize how lockdowns impact the contagion process and social welfare. Numerical analysis suggests that protracted, open-ended lockdowns are generally suboptimal, bringing into question the policy responses seen in many countries.

© 2021 Published by Elsevier B.V.

1. Introduction

The emergence of the SARS-CoV-2 virus and the onset of the COVID-19 pandemic motivated many governments to bring to a stand-still all human activity, social and economic, for many months. The stated objective is to slow down contagion and prevent healthcare systems from being overwhelmed. Many countries have gone to the extreme of imposing long-lasting and drastic “lockdowns” (e.g., China, Italy, Spain, UK), i.e., mandatory stay-at-home orders, business closures, and sweeping limitations to the freedom of movement. These lockdowns (also known as NPI's for nonpharmaceutical interventions) have been implemented on an open-ended basis, with a severity and duration primarily tied to the growth rate in infections. Apart from a few notable exceptions (e.g., Sweden), most countries have sought to minimize a single risk, that of contagion from the SARS-CoV-2 virus, without fully accounting for the economic and social consequences of doing so. In the aftermath of these interventions, a public debate emerged questioning the optimality of these policies. Are these drastic interventions optimal from a social welfare perspective? Should we keep them in place to address future epidemics? Should their implementation be open-ended?

[☆] The authors thank the Editor, Andrés Carvajal, as well as an anonymous associate editor and two reviewers for many helpful comments on a previous version of this paper.

* Corresponding author at: Economic Science Institute, Chapman University, One University dr., Orange, CA 92866, United States of America.

E-mail addresses: camera@chapman.edu (G. Camera), alessandro.giofré@unifi.it (A. Giorfré).

The answer to these questions partly depends on how one models the relevant economic tradeoffs. This paper develops an analytical framework that makes explicit the process of contagion, and ties it to the frequency of economic activity. The model economy has a constant population composed of individuals who can earn income only in periods in which they meet a trade partner. Meetings occur on a market where a matching process pairs individuals at random—all pairs generate a deterministic flow of income, and dissolve at the end of the period. The model assumes transmissibility via asymptomatic individuals, a central feature of the COVID-19 epidemic, as well as no cost from trading while asymptomatic. This implies that individuals who are unaware of being infected, have no incentive to stay out of the market and, hence, can spread the disease by meeting healthy trade partners. Repeating this random matching process period after period is how the epidemic spreads over time. It is assumed that precluding business activity by closing the market, i.e., imposing a lockdown, stops further contagion. Apart from the possibility of reaching herd immunity, no other intervention is assumed to exist to manage the progression of the disease. Therefore, lockdowns are the go-to policy to reduce healthcare costs, which are assumed proportional to the spread of the infection across the population. The severity of the policy intervention corresponds to the lockdown duration.

Our analysis is divided into two parts. First, we lay out the mathematical machinery needed to characterize the contagion process for general interventions, ranging from minimal as in Sweden to extreme as in China or Italy. The initial step is to construct transition matrices that determine the path of the

infection, for any initial state of the infection. These are then used to calculate the dynamic evolution of the epidemic when we impose lockdowns of various degree of severity. The analysis considers two scenarios, depending on whether the infection can or cannot die out by achieving herd immunity by medical means or naturally.

Second, we construct a measure of social welfare that combines individual payoffs from trade with expected healthcare costs associated with the spread of the disease. Lockdowns now delineate a tradeoff: more drastic interventions prevent overwhelming the healthcare system but destroy income flows. Numerical analysis suggests that welfare nonlinearly responds to the severity of the intervention, which leads to two main results. Imposing a lockdown is generally welfare-enhancing if the infection spreads easily. However, the welfare benefit rapidly dissipates as the lockdown length increases, and turns into a welfare loss eventually. If the infection is detected early and has reached only a small subset of the population, then imposing an extreme lockdown is counterproductive in terms of social welfare. Open-ended lockdowns are not necessarily optimal either, especially if the epidemic can be brought under control via herd immunity.

Intuitively, in our model the social gains from not overburdening the healthcare system are eventually overtaken by the economic losses stemming from further reductions in income flows. This is why extreme lockdowns are largely suboptimal. In fact, the analysis also reveals that naïvely matching the severity of the intervention to the spread of the infection is not the most logical policy because welfare gains are non-linear. Overall, this exercise suggests that policymakers should tread carefully. To the extent that healthcare conditions and income-generating processes are country-specific, the model indicates that the tradeoffs associated with lockdown policies are also country-specific. In other words, there is no “one-size-fits-all” kind of policy, which seems opposite to the adoption patterns seen so far, where many governments simply followed similar policies.

There is a voluminous literature on infectious diseases, an extensive review of which is beyond the scope of this paper. Due to space limitations, we refer the reader to the recent survey in Avery et al. (2020) and here we explain what our study adds to the existing literature. The novelty of our contribution lies in the approach to studying the diffusion of epidemics, which relies on the theory of random matching. This technique allows us to offer a framework that makes explicit the transmission of a disease in the population—in contrast to the standard epidemiological model, which uses a reduced-form approach. This framework is then used to assess the economic optimality of policy interventions based on the lockdowns imposed in the recent past.¹

¹ Our objective to study the welfare impact of health policies is shared by other recent works, all of which modify the standard SIR framework while maintaining its basic reduced-form approach. For example, Eichenbaum et al. (2020) assume that economic decisions affect the path of the disease because consuming and working less reduces the probability of becoming infected (hence, transition probabilities between health states). This gives rise to an externality that can be partially internalized by imposing limits to consumption and work activity (such as it happens in a lockdown). The model in Goenka et al. (2014) also considers feedback effects from disease to economic decisions. They embed the SIR model into a neo-classical growth model where investment in health capital alters the incidence of the disease, and the latter affects labor supply. They solve an optimal control problem, showing that both a disease-free steady state and a disease-endemic steady state may exist. The study in Alvarez et al. (2020) also follows the typical approach in the epidemiology literature where the evolution of the epidemic is a function of exogenous parameters, and extends it by embedding an optimal control problem, whereby a social planner chooses the diffusion parameter to maximize social welfare. An optimal control problem is also at the heart of Acemoglu et al. (2020), which extends the canonical single-group SIR model to a multi-group version with group-specific parameters. In particular, contact rates are governed by a matching function

To elaborate on this, start by noting that the typical model used in the epidemiology literature – known as the SIR model – is based on three possible states for an individual: S for susceptible (to infection), I for Infected, and R for Recovered. The evolution of these three mutually exclusive states is governed by laws of motion that underlie a reduced-form process of contagion. We retain the three-state representation typical in the literature, and innovate by constructing an *explicit* model of contagion, which is based on a pairwise random meeting process and the analysis technique developed in Camera and Giorfré (2014). In this manner, the model allows us to track the evolution of the disease when the markets are open or close – thanks to the explicit matching process – and thus study the optimality of lockdown policies under alternative scenarios for the initial state of the infection and the probabilistic nature of reaching immunity.

Given the persistent nature of the current epidemic, we also allow recovered individuals to be potential candidates for re-infection—something that is atypical in epidemiological models, where the recovered cannot be re-infected. To explain, in the standard SIR framework the number of susceptible individuals decreases over time due to the recovery process. In this case, threshold parameters exist such that an absorbing state is eventually reached where the infection vanishes as enough individuals contract the disease and herd immunity is naturally achieved.² By contrast, we work with a model where everyone in the population is generally susceptible to the disease until a point where the disease can be fully eradicated. It is assumed that in each period there is a probability that contagion stops, and until that happens everyone remains susceptible to infection, even recovered individuals. In this manner, a state of immunity is reached probabilistically and simultaneously by everyone in the population. This is a mathematically convenient way to capture a prominent aspect of current thinking behind lockdown policies: the COVID-19 disease is so dangerous and hard to contain that the epidemic *must* be stopped with a mass-vaccination campaign. This set-up allows us to trace a most favorable scenario for imposing lockdowns, which is when they delay the progression of the disease while medical and pharmacological interventions are being developed to address the problem.³

The paper proceeds as follows. Section 2 presents the model economy. Section 3 characterizes the dynamic process of contagion. Section 4 studies the impact of lockdowns on the spread of the epidemic in a baseline, worst-scenario model when there is no herd immunity threshold. In Section 4.4, we extend the analysis to a richer model where herd immunity can be reached by naturally acquired immunity and medical discovery. In Section 5, we apply this machinery to determine how interventions of varying severity impact social welfare; this is done by studying income losses and healthcare cost savings associated with lockdowns via numerical experiments. Broader policy implications of our analysis are discussed in Section 6, which concludes the study.

that is group-specific. In this context, policies that apply differential lockdowns across groups are superior to uniform policies that identically affect everyone in society. The main difference between these frameworks and ours is that we do not use a reduced-form approach; we construct an explicit meeting process that determines the spread of the infection.

² The theoretical analysis in Busenberg and van den Driessche (1990) shows how three threshold parameters govern population growth, the growth of cases, and the possibility that the disease becomes endemic.

³ “But there is one fact I want every American to know: People who are not fully vaccinated can still die every day from COVID-19” (President Joe Biden White House, 2021). This statement mirrors the official WHO policy that “Herd immunity against COVID-19 should be achieved by protecting people through vaccination, not by exposing them to the pathogen that causes the disease” (WHO, 2020b). This policy has led many countries to mandate lockdowns in order to gain the time necessary to set-up and execute mass-vaccination campaigns (e.g., see Wall Street Journal, 2021).

2. Modeling the economy

Time is discrete and infinite. The economy is composed of a constant population of $N = 2n \geq 4$ workers who can trade on a decentralized market. In every period $t = 0, 1, 2, \dots$, the market can be open or closed. If the market is closed, all individuals remain isolated and each individual obtains a payoff y . If the market is open, then individuals meet to trade in pairs. A policymaker chooses whether the market is closed or open in a period, which is discussed later.

Here, we note that if the market is open, then it is possible (but not certain) for an individual to meet a random trade partner. We interpret a meeting as a trading situation that is advantageous to both individuals. Considering a generic pair (i, j) , individual i obtains payoff $\bar{y} > 0$ from being in the meeting. Symmetrically, we have a payoff \bar{y} for individual j . The payoff corresponds to the instantaneous utility assumed in matching models of the labor market (Mortensen and Pissarides, 1994). We normalize $y = 0$ to underline that economic activity is beneficial, as it is necessary to create economic surplus. Market inactivity harms economic welfare because it does not allow surplus to be generated.

Individuals discount future payoffs with a common discount factor $\delta \in (0, 1)$. Letting $y_t \in \{y, \bar{y}\}$ denote the generic payoff to an individual at date t , the expected payoff to any individual at the start of the economy is therefore

$$\sum_{t=0}^{\infty} \delta^t y_t.$$

It is assumed that an infectious disease exists in the population, which can be transmitted only when individuals meet to trade. That is to say, in the model contagion occurs solely via business activity and not social interactions. This process is described in what follows.

The health status of individuals. Partition the population of workers N into three sets denoted *healthy*, *symptomatic* and *asymptomatic*. These last two sets are collectively called *infected* individuals. It is assumed that an asymptomatic individual becomes symptomatic at the start of a period with constant probability $s \in (0, 1)$. An infected individual can fully recover from a period to the next with probability $a \in (0, 1)$, thus regaining a healthy state and becoming non-infectious, a standard assumption in the epidemiology literature.⁴ This probability of recovery is independent of symptoms. As a result, the healthy set includes both individuals who were never infected or those who were and recovered. We thus have three possible states for an individual: *infected and symptomatic*, *infected and asymptomatic*, and *healthy* (never infected and recovered); see Fig. 1.

We make two assumptions that match the empirical characteristics of the COVID-19 disease and partly depart from the standard SIR model discussed in the epidemiology literature.

Assumption 1. An individual who is *recovered* is not infectious but is susceptible to future infection.

This assumption sets our model apart from the standard epidemiology literature, which typically assumes that recovered individuals cannot be re-infected and cannot infect others (e.g., see Avery et al., 2020).

Assumption 2. An individual who is *infected and asymptomatic* cannot be distinguished from a *healthy* individual.




Fig. 1. State transitions.

This implies that the only state that can be identified is being symptomatic, thus allowing for these individuals to be isolated. Neither the asymptomatic individual nor her counterparts can detect the presence of the infection. It follows that in our model trading activity can transmit contagion when asymptomatic and healthy individuals meet to trade. It is assumed that in such a meeting contagion occurs with probability $\eta \in (0, 1]$, independent of whether or not the healthy individual is someone who recovered from the disease in the past. In other words, it is assumed there is no permanent immunity to the disease and that it cannot be rooted out by naturally acquired immunity or medical means. That is, we make the infection very difficult to handle, so we give best shot at a lockdown policy to be welfare enhancing.⁵ Since not all infected individuals are present in the market – symptomatic individuals are excluded – to calculate the transmission rate of the disease we must first discuss how individuals are matched into potential trade meetings. This is done in the following subsection.

Trade meetings. Consider a period when the market is open. At the start of the period a matching process determines a *proposed* partition of the entire population of workers, into pairwise trade meetings. This means that every worker is assigned to a proposed pair, for the period. Pairs are selected using a uniform random matching process, e.g., as in Diamond (1982). Therefore, in each period t , the probability that individual i is assigned to meet any other individual is $\frac{1}{N-1}$. No meeting can last for more than one period, meaning that rematching takes place in each new trading period.

The population partition is only “proposed” because not all meetings necessarily take place: we assume that symptomatic workers are prevented from joining a meeting (e.g., they are sick, so cannot work). As a result, only healthy and asymptomatic workers join their proposed meeting. To clarify, consider Fig. 2, which displays the timeline of events in a period. First, meetings are proposed (a partition is proposed), then all workers join their assigned meeting but for symptomatic individuals. This is an analytically convenient way to maintain tractability, because we can run the matching process on a stationary population of size N .

⁴ We also assume that the state of an infected individual (symptomatic or not) is probabilistically determined at the start of each period. This allows us to avoid tracking the history of symptoms of individuals, which is necessary to keep the state space manageable and the model tractable.

⁵ We later relax this assumption by introducing the possibility of medical progress or naturally acquired immunity as a way to root out the disease.




Fig. 2. Timeline of events.

We now explain the remaining aspects of the timing of events in a period. If there are k infected individuals in the economy after meetings take place in a period t , then at the start of the following period they may decline to $\tilde{k} \leq k$ because some of them might recover before a new set of meetings takes place. At that point a set of $N/2$ trades is proposed by assigning everyone in the population to a trade pair, using a uniform random matching process. After this assignment, some infected individuals may show symptoms and are prevented from meeting their assigned trade partner. As a result, not all proposed trades take place. In the meetings that do take place, some will involve asymptomatic individuals who may end up transmitting the infection to their healthy partner. Hence, in $t+1$ after meetings take place we have $k' \geq \tilde{k}$ infected individuals. This process repeats itself indefinitely and governs the progression of the infection in the economy.

Now consider a proposed trade meeting. We say that we have a *mixed match* if the meeting involves an infected and a healthy individual. From an ex-ante perspective, a mixed match results in a new infection with probability

$$p := (1 - s)\eta.$$

This is the probability that the infected individual in the proposed mixed match remains asymptomatic in the period, and infects her healthy counterpart. With the complementary probability $1 - p$ the proposed meeting does not result in a new infection because (i) if the infected is symptomatic (with probability s), then she cannot enter the meeting and (ii) if the infected is asymptomatic, then contagion does not occur with probability $1 - \eta$.⁶

Summing up, in our model the infection is transmitted by asymptomatic individuals, who cannot be recognized and isolated, in contrast with those who show symptoms—who can be excluded from all trading activity. Hence, the asymptomatic are the main channel of contagion, which reflects the empirical observation that the SARS-CoV-2 virus is infectious even without symptoms. Two other empirically relevant features of the model are that the disease is not necessarily endemic in the population (recovery is possible), and past exposure to the infection is not a guarantee of permanent immunity. Finally, the empirical observation that the disease has a very low fatality rate in the working-age population motivated us to consider a zero fatality rate since we only consider contagion occurring through economic interactions.⁷ This is also analytically convenient as doing

so makes the population size stationary. We now proceed by showing how to trace the dynamics of the contagion in our model.

3. Characterizing the evolution of the disease

As seen above, the evolution of the disease in the economy involves two separate processes. The contagion process that occurs through trading activity, and the recovery process that takes place in-between separate rounds of trading. This section discusses these two processes, starting with the contagion process operating in the market, which is based on the study in Camera and Giofr  (2014).

3.1. Contagion through business activity

Consider the start of a period when the market is open. For notational convenience let $\tilde{k} = k$ denote the number of infected individuals. That is to say, there are $k = 1, \dots, N$ infected individuals who might interact in the market.

We start by deriving the probability $Q_{kk'}(N)$ that $k' \geq k$ individuals are infected by the end of a period of market activity. Two sources of randomness affect this probability: the matching process, which determines how many meetings occur between infected and healthy individuals, and the biological process, which determines if the infection is transmitted in these meetings.

Meetings between healthy and infected individuals. Here, we derive the probability $\lambda_{k\ell}(N)$ that, if we have k infected individuals and $N - k$ healthy individuals, then there will be $\ell = 0, 1, \dots, \min(k, N - k)$ pairs composed of one infected and one healthy individual. We call these pairs “mixed matches”, i.e., the only meetings where contagion can occur in the model. We have:

$$\lambda_{k\ell}(N) := \begin{cases} \frac{\ell! \binom{k}{\ell} \binom{N-k}{\ell} (k-\ell-1)!! (N-k-\ell-1)!!}{(N-1)!!} & \text{if } \ell \in L_k \\ 0 & \text{if } \ell \notin L_k, \end{cases} \quad (1)$$

where⁸

$$\ell \in L_k := \begin{cases} \{0, 2, 4, \dots, \min(k, N - k)\} & \text{if } k \text{ is even,} \\ \{1, 3, 5, \dots, \min(k, N - k)\} & \text{if } k \text{ is odd.} \end{cases} \quad (2)$$

To derive (1), notice that if k is even (odd), then the number ℓ of mixed matches cannot be odd (even), which explains why $\lambda_{k\ell}(N) = 0$ in (1) if $\ell \notin L_k$. Now consider $\ell \in L_k$. There are $\binom{k}{\ell}$

⁶ This is a convenient way to model the transmission process as opposed to tracking individual histories of those who have been exposed to the virus. This allows us to avoid tracking all individual histories and characterize the distribution of all individuals across a minimum of six different states, healthy and never exposed, exposed and asymptomatic, exposed and symptomatic, exposed and recovered, exposed and deceased, exposed and relapsed.

⁷ The Infection Fatality Rate (IFR) from COVID-19 disease for individuals in the working-age category is small. In the US, it is 0.02% for a 20–49 year old and 0.5% for a 50–69 year old. See CDC (Table 1 2020). Other studies report the average global IFR from COVID-19 as being 0.15% (Ioannidis, 2021a) and the median IFR for the population below 70 years of age as 0.15% (Ioannidis, 2021b).

⁸ In expression (1) we use the standard notation $!!$ of the double factorial. For an integer n , the double factorial is recursively defined as $n!! = n \cdot (n - 2)!!$ and, by definition, we have $0!! = 1$ and $(-1)!! = 1$.

possible ways to draw ℓ individuals from the set of those who are currently infected on the market (k individuals). Similarly, there are $\binom{N-k}{\ell}$ possible ways to draw ℓ individuals from the set of those who are currently healthy on the market ($N-k$). All healthy individuals, be they infected and recovered or never infected, are on the market trading. Hence, there are

$$\binom{k}{\ell} \binom{N-k}{\ell}$$

possible ways to draw ℓ infected and ℓ healthy individuals. Consider now all possible ways to form ℓ mixed matches. Fix an infected individual and match him to any of the ℓ healthy individuals. Once this match is formed, fix another infected individual and match him to any of the $\ell-1$ remaining healthy individuals. Repeating the process until everyone is matched, there are

$$\ell \cdot (\ell-1) \cdot (\ell-2) \cdots 3 \cdot 2 \cdot 1 = \ell!$$

possible mixed matches between ℓ infected and ℓ healthy individuals.

Hence, the number of pairings that give rise to at least ℓ mixed matches is

$$\ell! \binom{k}{\ell} \binom{N-k}{\ell}.$$

Since we are interested in finding the number of pairings that generate exactly ℓ mixed matches, we need to make sure that the pairings among those who are left do not generate additional mixed matches. In other words, we need to make sure that the remaining $k-\ell$ infected individuals are matched among themselves, and so are the remaining $N-k-\ell$ healthy individuals. Recall that $k-\ell$ and $N-k-\ell$ are necessarily even numbers. Considering the set of remaining infected individuals and fix one of them. Match him to one of the remaining $k-\ell-1$ infected individuals of this set. Once this match is formed, fix another infected individual and match him to one of the $k-\ell-3$ infected individuals who are left. Repeating this procedure until all $k-\ell$ infected individuals have been matched among themselves gives

$$(k-\ell-1) \cdot (k-\ell-3) \cdots 3 \cdot 1 = (k-\ell-1)!!$$

possible pairings. Similarly, we have $(N-k-\ell-1)!!$ possible pairings among the remaining healthy individuals.⁹

Therefore, the number of pairings that generate exactly ℓ mixed matches is

$$\ell! \binom{k}{\ell} \binom{N-k}{\ell} (k-\ell-1)!! (N-k-\ell-1)!!.$$

Finally, using the same argument above, the number of possible proposed pairings in the population of N individuals, is $(N-1)!!$. This concludes the derivation of the probability $\lambda_{k\ell}(N)$ of having ℓ market meetings among healthy and infected individuals, when there are k infected individuals.¹⁰

⁹ Note that if $\ell = k$ then by definition of the double factorial we have $(k-\ell-1)!! = 1$ and if $\ell = N-k$ then we have $(N-k-\ell-1)!! = 1$.

¹⁰ Note that $\lambda_{k\ell}(N)$ for $\ell \in L_k$ simplifies to

$$\lambda_{k\ell}(N) = 2^\ell \frac{\Gamma(k+1)\Gamma(N-k+1)\Gamma(\frac{N}{2}+1)}{\Gamma(\ell+1)\Gamma(\frac{k-\ell}{2}+1)\Gamma(\frac{N-k-\ell}{2}+1)\Gamma(N+1)}$$

where we have used $m! \equiv \Gamma(m+1)$, with Γ representing the gamma function, and, for $m = 2k-1$ (odd number), we have used $m!! \equiv \frac{(2k)!}{2^k k!}$. This functional transformation of $\lambda_{k\ell}(N)$ simplifies the calculations in Matlab, where we can use the `gamma` function to work with log linearized expressions.

Transmission of the infection conditional on a mixed match. The infection might not spread in all mixed matches because we assumed that transmission of the disease between a healthy and an infected individual occurs with probability $p < 1$. This has two implications. First, the number ℓ of mixed matches that can result in $k' - k \geq 0$ new infections is $\ell = k' - k, \dots, \min(k, N-k)$. Second, the probability that we have enough market meetings capable of creating $k' - k$ new infections is

$$\sum_{\ell=k'-k}^{\min(k, N-k)} \lambda_{k\ell}(N),$$

i.e., the probability that we have at least $k' - k$ mixed matches. By construction, the sum of probabilities $\sum_{\ell=0}^{\min(k, N-k)} \lambda_{k\ell}(N) = 1$.

To derive the conditional probability that transmission occurs in exactly $k' - k$ of these $\ell \geq k' - k$ matches start by considering the case $\ell = k' - k$. Here, the probability that the infection is transmitted in all mixed matches is $p^{k'-k}$. Instead, if $\ell > k' - k$, then transmission occurs in exactly $k' - k$ cases with probability

$$\binom{\ell}{k'-k} p^{k'-k} (1-p)^{\ell-(k'-k)}$$

Putting together all this information leads to the following:

Lemma 1. *The probability that $k' \geq k$ individuals are infected by the end of a period of market activity is:*

$$Q_{kk'}(N) = \sum_{\ell=k'-k}^{\min(k, N-k)} \lambda_{k\ell}(N) \binom{\ell}{k'-k} p^{k'-k} (1-p)^{\ell-(k'-k)}. \quad (3)$$

In deriving $Q_{kk'}(N)$ we have only considered states in which someone is infected, $k, k' > 0$. For the remaining cases $k = 0$ and $k' = 0$ we define the transition probabilities $Q_{00} = 1$, $Q_{0k'} = 0$ for all $k' > 0$, and $Q_{k0} = 0$ for all $k > 0$. The first two definitions imply that the zero-infection state is absorbing (if the disease is non-existent we cannot have new infections) and the third simply follows from the observation that market interaction can only generate new infections.

Note that the random process determining the number of infected individuals is a finite Markov chain because the transition probability $Q_{kk'}$ depends only on the state $k = 0, 1, \dots, N$ at the start of a trading period and not on the entire history of infections.¹¹

These considerations lead to the following.

Corollary 1. *Let the market be open at the start of a period. The spread of the infection during the period is fully described by the $(N+1) \times (N+1)$ upper-triangular Markov matrix \mathcal{Q} :*

$$\mathcal{Q} := \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & Q_{11} & Q_{12} & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & Q_{22} & Q_{23} & Q_{24} & 0 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & Q_{N-2,N-2} & Q_{N-2,N-1} & Q_{N-2,N} \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & Q_{N-1,N-1} & Q_{N-1,N} \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 1 \end{pmatrix}.$$

¹¹ The proof of Lemma 1 immediately implies that all $Q_{kk'}(N) \leq 1$, as by construction they are probabilities. Although the analytical calculation $\sum_{k'=0}^N Q_{kk'}(N) = 1$ is not straightforward for a generic N , we note that this holds because $Q_{kk'}(N) = 0$ for $k' < k$ and $k' > \min(2k, N)$, and the states with positive transition probability $k' = k, \dots, \min(2k, N)$ are mutually exclusive. By means of example, if $N = 4$, $p = 1/2$, $k = 2$ we have $\lambda_{20} = 1/3$, $\lambda_{21} = 0$, $\lambda_{22} = 2/3$, and $Q_{22}(4) = \lambda_{20} + \lambda_{22}(\frac{1}{2})^2 = \frac{1}{3} + \frac{1}{6}$, $Q_{23}(4) = \lambda_{22}(\frac{1}{2})^2 = \frac{1}{3}$, $Q_{24}(4) = \lambda_{22}(\frac{1}{2})^2 = \frac{1}{6}$, so $\sum_{k'=2}^4 Q_{2k'}(4) = 1$.




Fig. 3. An open-ended lockdown intervention (T, j) . **Notes:** The lockdown policy (T, j) comes into effect in τ . The market reopens regularly every T periods, in $\tau + T, \tau + 2T + j, \dots$. O = market is open.

Because the number of infections is a finite Markov chain the square matrix \mathcal{Q} is a Markov matrix. An important property of Markov matrices is that the transition probabilities taking place after $j \geq 1$ steps is determined by the product \mathcal{Q}^j . Matrix \mathcal{Q} is a central element to calculate the evolution of the disease in the economy, but not the only one. The reason is that $Q_{kk'}(N)$ only tells us the probability that $k' - k \geq 0$ new infections occur as a consequence of market activity. It does not account for the possibility of recovery from the disease, in which case the infection might decline or even be completely eradicated. This additional component is discussed in what follows.

3.2. A random recovery process

Fig. 2 indicates that the number of infected individuals present at the start of a trading period depends on the recovery process that occurs between the end of a period and the beginning of the next. This process is next described.

Assume that an individual who results in being infected by the end of a period, recovers with probability a by the beginning of the following period. This implies that if we start with k infected individuals, and some recover, then at the beginning of next period we may have $\tilde{k} \leq k$ infected individuals.

We use $R_{k\tilde{k}}$ to denote the probability that, given $k = 0, 1, \dots, N$ individuals being infected at the end of a period, $k - \tilde{k}$ have recovered by the start of the following period, where

$$R_{k\tilde{k}} = \binom{k}{k-\tilde{k}} a^{k-\tilde{k}} (1-a)^{\tilde{k}} \quad \tilde{k} = 0, \dots, k. \quad (4)$$

It should be clear that if $k = 0$, then the contagious process stops and the disease is permanently eradicated. So, given $k = 0, 1, \dots, N$ infected individuals, the disease is eradicated with probability $R_{k0} = a^k$. With the complementary probability, it continues.

Again, the random process determining the number of infected individuals between the end of a period t at the beginning of $t + 1$ is a finite Markov chain because the transition probability $R_{k\tilde{k}}$ depends only on the state $k = 0, 1, \dots, N$, i.e., the number of infected individuals at the end of the period (not on the entire history of infections).¹²

This discussion immediately implies the following.

Lemma 2. Let there be $k = 0, 1, \dots, N$ infected individuals at the end of a period. The number of infected individuals at the start of the following period is fully described by the $(N + 1) \times (N + 1)$

lower-triangular Markov matrix

$$\mathcal{R} := \begin{pmatrix} 1 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ R_{10} & R_{11} & 0 & 0 & \dots & 0 & 0 & 0 \\ R_{20} & R_{21} & R_{22} & 0 & \dots & 0 & 0 & 0 \\ R_{30} & R_{31} & R_{32} & R_{33} & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \vdots \\ R_{N-2,0} & R_{N-2,1} & R_{N-2,2} & R_{N-2,3} & \dots & R_{N-2,N-2} & 0 & 0 \\ R_{N-1,0} & R_{N-1,1} & R_{N-1,2} & R_{N-1,3} & \dots & R_{N-1,N-2} & R_{N-1,N-1} & 0 \\ R_{N,0} & R_{N,1} & R_{N,2} & R_{N,3} & \dots & R_{N,N-2} & R_{N,N-1} & R_{N,N} \end{pmatrix}.$$

In the next section, we show how to use matrices \mathcal{R} and \mathcal{Q} to calculate the evolution of the number of infected individuals over time.

4. How lockdowns affect the contagious process

In each period the market can be either open or closed. If the market is open, then the random matching process proposed $N/2$ trade meetings. If the market is closed, then everyone remains unmatched for the period.

4.1. The “lockdown” policy intervention

Market interventions are imposed by an external authority (a government), which can select to close the market for any desired extended length of time, without restrictions and without consultation with the population. This is what we call a “policy intervention”.

The policy is completely described by the two parameters (T, j) , with $T, j \geq 0$. The policy specifies an initial *lockdown phase* composed of T consecutive rounds of market inactivity, followed by a *reopening phase* consisting of j consecutive trading rounds. This means that the policy’s horizon is $T + j$ periods, after which the policy expires and a new policy can be considered. Fixing the policy horizon $T + j$ to some arbitrary value, the parameter T defines the severity of the lockdown; $T = 0$ corresponds to no lockdown. As T increases, the lockdown is stricter.¹³

As an illustration, consider a policy (T, j) that is repeated until the infection is eradicated. In this case we have what we call a *lockdown cycle* (T, j) . This kind of open-ended policy intervention is illustrated in Fig. 3. The policy comes into effect in period τ , i.e., the market is open up until period $\tau - 1$. All trading stops for T periods at regular intervals τ then $\tau + T + j$, and so on. The market is also (re)opened for j periods at regular intervals $\tau + T$, then $\tau + 2T + j$ and so on.

¹² By construction, $R_{k\tilde{k}}(N) \leq 1$ and $\sum_{\tilde{k}=0}^N R_{k\tilde{k}}(N) = 1$.

¹³ Alternatively, one can interpret $T/(T+j)$ as controlling what fraction of the entire set of market activities is forced to shut down until there is no more risk of infection. For example, if $T = 5$ and $j = 20$ then only 20% of market activities remain open ($= 1/5$), while 80% of the market is shut down. These are discrete jumps in market inactivity. The model can be generalized to attain smaller and more progressive reduction in market activity.

As we assumed that infection can only occur in trade meetings, the contagion spreads at random across the population only when markets are open. Whenever a lockdown comes into effect, trade stops for T consecutive periods, preventing further contagion and allowing some recoveries to occur. Therefore, the two Markov matrices \mathcal{Q} and \mathcal{R} allow us to characterize the evolution of the infection in the economy for any policy intervention.

With this machinery, we can calculate the expected number of infections for any given lockdown policy (T, j) and, its welfare consequences.

4.2. The expected number of infections

It is convenient to first consider the case $T = 0$, i.e., when there is no lockdown. Here, the contagion process in a period is governed by the matrix product $\mathcal{Q}\mathcal{R}$. This product accounts for the trading and recovery processes described in Fig. 2. The product $\mathcal{Q}\mathcal{R}$ is a stochastic matrix. It represents the Markov process that governs the law of motion of the number of infected individuals in the economy in any period t in which the market is open. For any given number $k = 0, 1, \dots, N$ of individuals that result infected at the start of period t , it gives us the probabilities that we will have $k' = 0, 1, \dots, N$ infected individuals at the start of period $t + 1$. This product accounts for two components affecting the transition probabilities: \mathcal{Q} determines the new infections that can emerge as a consequence of trading activity in period t ; this is an intermediate state reached in period t . The final state is determined by \mathcal{R} , which gives us the probabilities that – after trading takes place – some infected individuals will recover by the start of $t + 1$.

Because the product $\mathcal{Q}\mathcal{R}$ is a stochastic matrix of the Markov type, we can iterate on it to calculate the number of new infections expected after any given number $j \geq 1$ of consecutive trading periods. In this case $(\mathcal{Q}\mathcal{R})^j$ governs the law of motion of the number of infected individuals in the economy between the start of period t and the start of period $t + j$.

Based on the above, let us now consider an economy of size N where there are $k = 0, 1, \dots, N$ infected individuals at the start of a period. The expected number of infected individuals after $j \geq 1$ consecutive trading periods is

$$\mu_k(j) := e_k^\top (\mathcal{Q}\mathcal{R})^j \kappa.$$

Here, $\kappa = (0, 1, \dots, N)$ is a column vector that contains all possible numbers of infected individuals in the economy (including the 0 absorbing state, when the disease is eradicated). The vector e_k is the $(N + 1)$ -dimensional column vector with 1 in the $(1 + k)^{\text{th}}$ position and 0 everywhere else. The transpose of this vector, denoted e_k^\top , selects the $(1 + k)^{\text{th}}$ row of matrix $(\mathcal{Q}\mathcal{R})^j$, i.e., the state of the economy corresponding to $k \geq 0$ infected individuals. The non-zero elements of that row are the probabilities to transition from k to $k' = 0, 1, \dots, \min(2k, N)$ infected people by the end of j consecutive trading periods.

Now consider the case when $T \geq 1$, i.e., when trading activity restarts after T periods of complete isolation of all traders. During the lockdown phase there cannot be further contagion and there can be some recoveries. Hence, \mathcal{R}^T determines the decline in the number of infections that we can expect from T rounds of lockdown, which is simply calculated by iterating T times transition matrix \mathcal{R} . When the lockdown is lifted the matrix product $\mathcal{Q}\mathcal{R}$ determines the spread of the infection in each trading period. Hence, we use $(\mathcal{Q}\mathcal{R})^j$ to determine the evolution of the infections if the market remains open for j periods, calculated by iterating j times transition matrix $\mathcal{Q}\mathcal{R}$.

With this machinery we can easily determine not only how a lockdown can slow down the evolution of infections, but also how reopening markets can speed up contagion. To explain, the

number of infections at the end of a T -period lockdown phase that is imposed after j periods of consecutive market activity is calculated using the transition matrix $(\mathcal{Q}\mathcal{R})^j \mathcal{R}^T$. Conversely, the transition matrix $\mathcal{R}^T (\mathcal{Q}\mathcal{R})^j$ allows us to determine the number of infections at the end of j consecutive periods of market activity after T periods of lockdown have been imposed. If we study policies by considering the end of market activity phase as the time reference, we thus have the following:

Lemma 3. Consider a lockdown policy (T, j) . If we start with $k = 0, 1, \dots, N$ infected individuals, the expected number of infected individuals at the end of the reopening phase following T periods of lockdown, is

$$\mu_k(T, j) = e_k^\top \mathcal{R}^T (\mathcal{Q}\mathcal{R})^j \kappa.$$

The lemma immediately follows from direct calculation. To see how lockdowns affect the spread of the infection an example may be helpful.

Example: no lockdown vs. 1-period lockdown. Suppose that $k = 2$ persons are infected by the end of a period. Let $p = 0.1$, $a = 1/4$, and $N = 1000$. We wish to calculate the number of infected individuals after the first period of market activity for two cases: (i) no lockdown, in which case the market is open in the following period, and (ii) lockdown, in which case the market is closed for one period and reopens the period after the next. In each case we have five possible outcomes: the contagion completely stops by the end of the following period, or the number of infections declines to 1, remains at 2, increases to 3, increases to 4. Each of these events has an associated probability. Since

$$R_{20} = a^2 = 0.0625, \quad R_{21} = 2a(1 - a) = 0.375,$$

$$R_{22} = (1 - a)^2 = 0.5625,$$

then, contagion stops with probability $R_{20} \approx 0.063$ (no lockdown) and $R_{20}R_{00} + R_{21}R_{10} + R_{22}R_{20} \approx 0.191$ (1-period lockdown). The number of infections declines to 1 with probability $R_{21}Q_{11} \approx 0.337$ (no lockdown) and $R_{21}R_{11}Q_{11} + R_{22}R_{21}Q_{11} \approx 0.443$ (1-period lockdown). The number of infections remains at 2 with probability $R_{21}Q_{12} + R_{22}Q_{22} \approx 0.493$ (no lockdown) and $R_{21}R_{11}Q_{12} + (R_{22})^2Q_{22} \approx 0.306$ (1-period lockdown). The number of infections grows to 3 with probability $R_{22}Q_{23} \approx 0.101$ (no lockdown) and $(R_{22})^2Q_{23} \approx 0.057$ (1-period lockdown). Finally, the number of infections grows to 4 with probability $R_{22}Q_{24} \approx 0.006$ (no lockdown) and $(R_{22})^2Q_{24} \approx 0.003$ (1-period lockdown).

As a result, the expected number of infections after the first period of market activity is 1.650, if the market is open tomorrow because there is no lockdown, and 1.238 if the market is closed tomorrow and reopens the period after.

4.3. How lockdowns affect the spread of the epidemic

We have now all the needed machinery in place to show how the duration of the lockdown, T , influences the spread of the epidemic in the basic model. In particular, we wish to determine the effectiveness of lockdowns in speeding up the attainment of a zero-prevalence state by leveraging the process of natural recoveries while preventing further contagion. In a follow-up section (Section 4.4), we extend this analysis to a richer model, in which the zero-prevalence state can be attained also via naturally acquired immunity or experimental medical procedures.

Numerical illustration. We start by showing that the probability of transmission is central to determine whether or not lockdowns are needed to bring contagion under quick control. To illustrate this point, consider Fig. 4, which shows three different scenarios for the transmissibility of the disease: low, medium, and high.




Fig. 4. Infection progression without intervention (herd immunity impossible). **Notes:** The population size is $N = 1000$, $a = 0.25$, $s = 0.2$ and there is no possibility of herd immunity. The three solid lines consider an initial infection rate of 1%, the three dashed lines an infection rate of 10%. For each initial infection rate, the three curves are drawn for “low”, “medium” and “high” transmissibility of the disease, corresponding to $\eta = 1/8, 3/8, 2/3$.

The illustration considers the baseline case when herd immunity is impossible, $N = 1000$, a recovery rate of $a = 1/4$ and $s = 0.2$. We vary the transmissibility of the disease in a trade meeting by varying η .¹⁴ For illustrative convenience, we trace the policy over a horizon of 100 periods. If we think of a period as a week, then we have approximately a two-year horizon, which allows us to contrast the short run and long run impact of policy implementation.¹⁵

In the figure, the share of infected individuals progresses according to the process of market interactions. For low and medium transmission rates, the natural recovery process prevents an increase in infections. That is to say, if the natural process of recovery is robust and the transmission rate is sufficiently low, then the infection rate does not get out of control and, even if markets remain fully open, it falls over time. Instead, the disease rapidly spreads in the population if the transmission rate is sufficiently high and does not decline. This brief illustration suggests that lockdowns play the most significant role as possible tools to curb contagion when the infection is highly transmissible, which is the case we focus on in studying how lockdowns affect the spread of the disease.

Consider a lockdown policy (T, j) that has two phases: a *lockdown phase* lasting T periods, when markets are closed and there is no contagion (only recoveries); and a *reopening phase* lasting j periods, when markets reopen and contagion restarts. We juxtapose a moderate to severe policy, which are differentiated by the length of the lockdown phase $T = 10, 20$, respectively. As T increases, the recovery process lasts longer, causing a greater reduction in the number of infected individuals before markets reopen.

¹⁴ According to WHO (2020a), 80% of COVID-19 infections are mild or asymptomatic, which motivates $s = 0.2$. A recovery rate of $1/4$ pins down a duration of infection of about four weeks (the inverse of a). There is no definitive way to pin down η , as the number of secondary infections generated from one infected individual varies greatly according to context and time (WHO, 2020a). The value $1/\eta$ pins down the expected number of meetings it takes an infected individual to spread contagion to one healthy trade partner.

¹⁵ A period does not have an absolute interpretation in our model (day, week, month) unless one specifies a discount factor. A discount factor that pins down the length of time encompassed by a period: the higher the value the smaller the length of time. This factor will become important when we calculate the welfare consequences of a lockdown policy in a later section.

For each of these T values, we consider two alternative durations j of the reopening phase: predetermined and state-dependent. In the predetermined case, after the lockdown ends markets remain open for j periods. In the state-dependent case, when markets reopen a lockdown is immediately re-imposed if the infection rate climbs above a pre-specified trigger level—hence j varies depending on the progression of the infection.

In the numerical experiments, we consider the case of a highly transmissible disease. The economy starts in period 0 with open markets and two initial infection rates, 1% and 10% ($k = 10, 100$ respectively, given $N = 1000$). Two different kinds of responses are considered: quick and delayed. Under a *quick response*, there is a short, 4-period interval before the start of the lockdown phase. Under a *delayed response*, the delay grows to 12 periods. We report the expected number of infections for moderate and severe lockdowns, in the top and bottom panels of Fig. 5, respectively. The dots pin down the period when lockdowns are imposed and the squares the period when markets are reopened, after which the reopening phase continues uninterrupted. The horizontal lines correspond to the initial infection rate; they also allow us to illustrate state-dependent policies according to which a new lockdown is imposed based on, alternatively, a high and a low infection trigger (10% and 1%). There are three main observations about these economies where herd immunity is assumed impossible.

Observation 1. *If the infection can be transmitted through economic activity, then a lockdown slows down the progression of the infection in proportion to its duration.*

Consider the top panel. For any initial infection rate, high or low, a moderate lockdown lowers the number of infections in the short run, but not in the long run. The expected share of infected individuals rapidly falls as soon as the lockdown is imposed, due to natural recoveries and the absence of additional contagion. This beneficial effect of the lockdown intervention is of short-duration. Infections quickly climb back up as soon as markets reopen, quickly surpassing the initial infection levels. By contrast, in the bottom panel the lockdown lasts twice as long. This lowers the long-run expected infection rate well below the initial levels, even for the high 10% initial rate. Intuitively, the longer duration of the intervention greatly reduces the number of infected individuals so the progression of the contagion is slow when markets reopen. Seen this way, more severe lockdowns more effectively reduce the expected number of infections in the long run. However, even in this case infection rates eventually climb back up. This last consideration leads to our second observation.

Observation 2. *If herd immunity is impossible, then lockdowns are ineffective at containing the epidemic in the long run, unless they are repeated and persistently restrictive.*

To illustrate this, consider the top panel of Fig. 5. Infections climb back up very quickly when markets reopen, and rapidly exceed the initial infection rate. In this sense, a one-time moderate lockdown is ineffective at containing the epidemic in the long run. To ensure low infection rates, the lockdown phase must be repeated. To illustrate this suppose a new round of lockdowns is triggered whenever the infection rate climbs up to a pre-determined threshold, either high (10%) or low (1%). For a low threshold, this occurs soon after markets reopen. For a high threshold the delay between the end of a lockdown phase and a new round of lockdowns is longer.

Now consider the bottom panel of Fig. 5. The severe lockdown more greatly reduces infections. Hence, if policymakers quickly impose a lockdown and intend to maintain the infection rate below, say, 10%, then there is no need to impose another round of

lockdowns. However, lockdowns must be re-imposed if the target infection rate is lower (say, 1%) or if lockdowns are imposed with delay. This suggests one more observation.

Observation 3. *Delays in policy implementation affect the long-run path and the need for repeated interventions.*

In each panel of Fig. 5, infections quickly decline when lockdowns are implemented. The faster markets are shut down the sooner infection rates drop. However, implementation delays may also have an impact in the long run as the infections expected after a quick and a delayed response take two different paths in the long run. The speed of policy response is clearly relevant for short-run management of the disease, which may be critical if there are tight constraints on healthcare resources in the short run. In addition, the illustration suggests that the speed of the response may also be relevant in the long run, especially for determining if lockdowns must be repeatedly applied to contain the disease.

Overall, these numerical illustrations suggest that – presuming that herd immunity is impossible – imposing a one-time lockdown cannot adequately curb the infection, even if the duration of the lockdown is quite long. This brings into question the emphasis on imposing one-time extreme lockdowns to contain the epidemic—the “flattening the curve” notion made popular in the media at the onset of the COVID-19 epidemic. The reason is that, since immunity is assumed impossible, as soon as markets reopen the disease starts to spread again. In this case, markets must be shut down as long as necessary for the process of natural recoveries to bring down the number of infections to zero. At that point, markets can be reopened. The alternative to such and extreme intervention is to repeat the lockdown to maintain infection rates at a manageable level. The open question is whether these conclusions change when herd immunity is possible, which is what we study next.

4.4. Generalization: herd immunity

In this section we extend the model to include the possibility of achieving herd immunity. For convenience, assume that in each period there is a time-invariant probability h that contagion stops. When this occurs, the number of new infections becomes zero forever. This can be thought of as the consequence of the mass-vaccination campaigns currently implemented in many countries (White House, 2021), or scientific discovery leading to the emergence of effective pharmacological interventions (e.g., the repurposing of ivermectin as recently reported in Kory et al., 2021). It can be also seen as a way of modeling naturally acquired immunity in reduced-form, i.e., without making explicit a process of gradual pathogenic exposure that subsequently gives rise to natural immunity.

The probabilistic herd-immunity assumption implies that we expect the disease will be eradicated $1/h$ periods in the future, since

$$h + 2(1-h)h + 3(1-h)^2h + \dots = 1/h.$$

Hence, $1 - (1-h)^n$ is the probability that herd immunity is achieved after n periods. Since this probability is independent of other model parameters, we can calculate an expression equivalent to $\mu_k(T, j)$, i.e., the number of infected individuals expected when a lockdown phase of T periods is followed by a reopening phase of j periods.

Lemma 4. *Let there be $k = 0, 1, \dots, N$ infected individuals the period before a lockdown policy (T, j) is imposed. Let $h > 0$.*

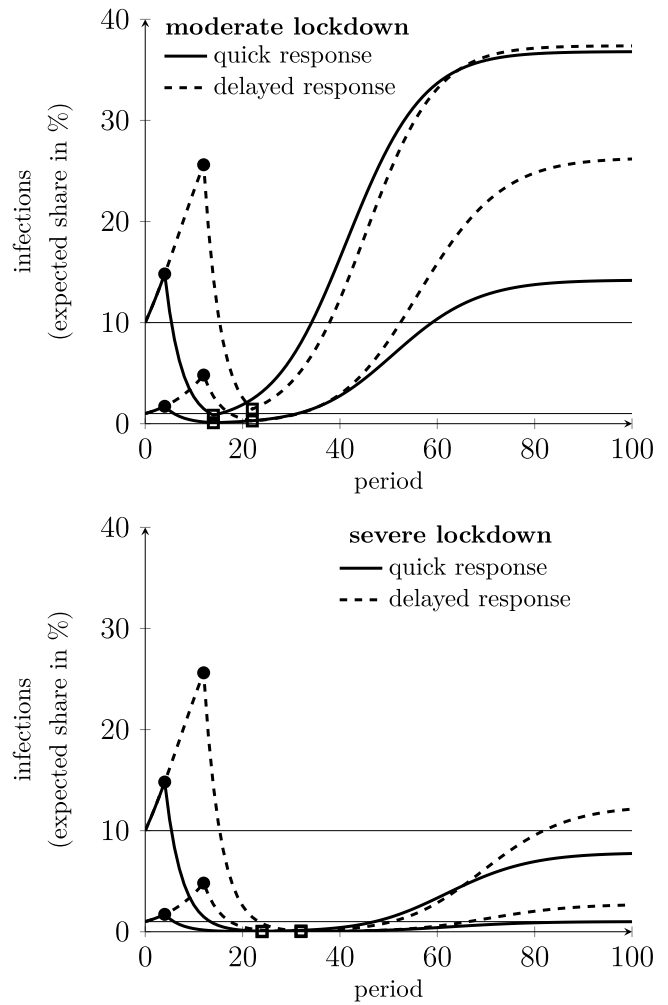


Fig. 5. Infection progression with intervention (herd immunity impossible). **Notes:** The population size is $N = 1000$, $a = 0.25$, $s = 0.2$ and there is no possibility of herd immunity. The top (bottom) panel considers a lockdown $T = 10$ ($T = 20$); these policies correspond to “moderate” and “severe” interventions. Each panel considers two possible initial infection rates, 1% and 10%, identified by the horizontal lines. The solid curves correspond to a “quick response” (the lockdown policy is implemented after 4 periods), while the dashed curves identify a “delayed response” (12 periods delay). Circles identify periods when the lockdown starts, the squares identify the start of the reopening phase.

The expected number of infected individuals at the end of the first lockdown policy is

$$\hat{\mu}_k(T, j) = (1-h)^{T+j} \mu_k(T, j).$$

The expected number of infected individuals after j consecutive rounds of trading is thus $\hat{\mu}_k(0, j) = (1-h)^j \mu_k(0, j)$.

Proof of Lemma 4. To prove the Lemma, consider a policy consisting of T rounds of lockdown followed by j trading rounds. At that point the policy expires and a (possible) new policy comes into effect. The probability that herd immunity is achieved at any point during these $T + j$ periods is thus:

$$h + (1-h)h + (1-h)^2h + \dots + (1-h)^{T+j-1}h \equiv 1 - (1-h)^{T+j}.$$

With the complementary probability $(1-h)^{T+j}$, herd immunity is not achieved. Given that this probability is independent of other factors in the model, the lemma is easily obtained by construction, using the result in Lemma 3. \square




Fig. 6. Infection progression without intervention (herd immunity possible). **Notes:** The population size is $N = 1000$, $a = 0.25$, $s = 0.2$ and $h = 0.01$. For other details see notes to Fig. 4.

Numerical illustration. Fig. 6 illustrates the expected progression of the infection without any kind of intervention. Here, $h = 0.01$, so if we interpret a period as a week, herd immunity is anticipated to be achieved in about two years.

For low and medium transmission rates, the expected path of the infection is similar to the case illustrated in Fig. 4. Intuitively, in those cases the natural recovery process is enough to curb contagion so adding the possibility of herd immunity does not change the overall picture. Instead, the possibility of attaining herd immunity is of primary importance when the disease is highly transmissible. In that case the progression of the disease is much more contained as compared to the case illustrated in Fig. 4. Compared to the no-herd-immunity case, the infection is expected to peak in the short run, and to reach a lower rate (about 10 percentage points less in the illustration). The problem is that the infection rate declines slowly after it reaches the peak. This suggests that one-time lockdowns may be helpful to manage the disease in the high-transmissibility scenario, because now they can complement the process underlying the development of herd immunity. This is illustrated in Fig. 7, which considers the same interventions studied earlier for the case $h = 0.01$ and high-transmissibility. Two main observations can be made.

Observation 4. *If herd immunity can be achieved, then repeated lockdowns may be unnecessary to bring the epidemic under control.*

The top panel in Fig. 7 shows that a one-time moderate intervention significantly alters the long-run path of the infection relative to the no-intervention case. However, the moderate intervention illustrated in the top panel does not succeed in bringing the long-run infection rate below the initial level. The bottom panel shows that this can be accomplished by imposing a more severe lockdown. Even for high initial rates, the one-time severe lockdown is expected to significantly reduce long-run infection rates, and to eventually bring them to zero.

Note that after reaching its peak, the expected share of the infected population eventually declines to zero in all cases considered. This decline occurs even if the lockdown is not repeated, i.e., if markets remain indefinitely open after period 100 in the figure. In this case, the policymaker can select a lockdown length T to contain the epidemic below a target rate, or to accelerate the convergence to a no-infection state. A new lockdown can further speed up the convergence to a no-infection state, and can do so

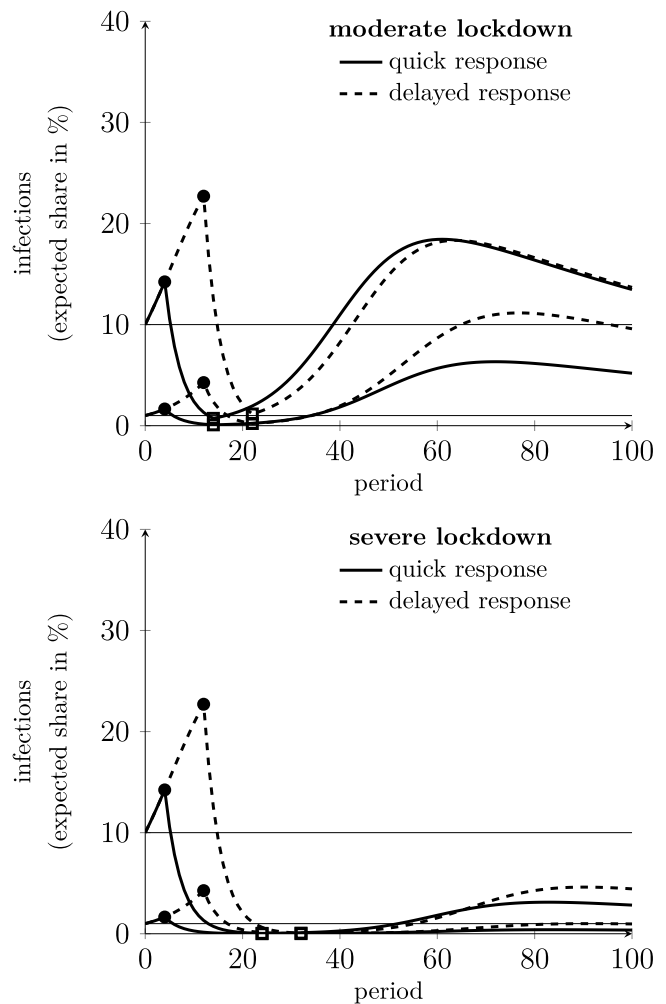


Fig. 7. Infection progression with intervention (herd immunity possible). **Notes:** The population size is $N = 1000$, $a = 0.25$, $s = 0.2$ and $h = 0.01$. For other details see the note to Fig. 5.

even if this subsequent intervention is more moderate than the first one (shorter duration).

In this sense, our analysis suggests that the repeated, extreme lockdowns that we have been experiencing do not appear to be generally necessary to contain the disease within manageable levels if herd immunity is possible. In this case, a one-time intervention can be sufficient to “flatten the curve” and contain the infection rate below a desirable target. We can make one more observation.

Observation 5. *If the epidemic is in its initial stages, then severe lockdowns are unnecessary to contain it, if the intervention is not delayed.*

This observation emerges from considering the top panel in Fig. 7 for a starting infection rate of 1%. Suppose that we want to maintain the infection rate below 10% so that new lockdowns would be imposed only if this level is reached when markets reopen. A moderate lockdown implemented with a short delay (solid line) quickly reduces the infection rate and once markets reopen the expected infection rate remains below 10% and eventually declines. However, implementing the lockdown with considerable delay (dashed line) fails to keep the infection rate below 10% when markets reopen. As the delay affects the long-run path of the infection, this requires the lockdown to be repeated in

order to achieve the desired target. The top panel also shows that if the infection is not in the initial stages, say we are at 10% and not 1%, then the moderate lockdown would not help managing the diseases in the long run, and the 10% target is rapidly overcome even if the intervention is implemented relatively quickly. In this case, we need a severe lockdown to bring the disease under control, maintaining the rate of infections below 10% (bottom panel).

We emphasize that these observations are simply illustrative, given the nature of the exercise conducted and the constraints imposed on the model. Yet, the numerical experiments provide some useful insights. The greater the severity of the lockdown intervention, the lower the share of the population expected to contract the disease. The impact is non-linear; a severe intervention is helpful to contain the epidemic, especially when the disease is already quite widespread. In this case, repetition of the intervention might be unnecessary to maintain control over the infection. This suggests that severe lockdowns can be useful to quickly bring a highly transmissible disease under control, but at the same time repeating the lockdowns does not seem to be generally necessary. The fundamental question is whether or not society benefits from this.

In order to look into this question, consider that lockdowns destroy income by precluding business activities. On the other hand, lockdowns prevent the healthcare system from being overwhelmed in the initial stages of an epidemic. They can also stave-off contagion long enough to ensure that – as the sick recover – healthcare capacity constraints do not bind in the long run. These two opposite economic considerations give rise to a tradeoff that we study in what follows.

5. Are lockdowns socially optimal?

In this section, we offer a measure of social welfare corresponding to per-capita expected payoffs under a lockdown policy (T, j) . The social welfare measure discussed in this section only considers the trade-off between healthcare costs and lost incomes during the periods that encompass the policy duration, which is a limitation introduced to maintain the analysis tractable. Another assumption made for tractability is that the effects of lockdowns do not linger in the long run; once lockdowns are lifted, markets immediately resume their normal operations.

Expected per-capita income from trading. We start by determining the maximum and minimum per capita income. In a representative period, no income is produced by those who cannot trade and $\bar{y} > 0$ for those who trade. Hence, an individual who trades in every period has lifetime payoff

$$v_0 := \frac{\bar{y}}{1 - \delta}.$$

Now consider the most extreme lockdown policy: markets do not reopen until no one is infected. Here, the individual earns nothing until the epidemic is brought under control, at which point the individual switches to trade forever after, earning v_0 . We have:

Lemma 5. *Let $k = 0, 1, \dots, N$ denote the number of infected individuals at the start of the most extreme lockdown policy. The expected payoff to an individual is*

$$v_k = e_k^T [\mathcal{I} - (1 - h)\delta\mathcal{R}]^{-1} h\mathbf{v}_0, \quad \mathbf{v}_0 := (v_0, v_0, \dots, v_0)^T.$$

Proof of Lemma 5. Suppose there are a number of $k = 0, 1, \dots, N$ infected individuals at the start of the period. Define the $(N + 1)$ -dimensional column vectors

$$\mathbf{v} := (v_0, v_1, \dots, v_N)^T \quad \text{and} \quad \mathbf{v}_0 := (v_0, v_0, \dots, v_0)^T.$$

These vectors define expected payoffs v_k given the number k of infected individuals at the start of the period, for the extreme case when the lockdown is only lifted when the disease is eradicated. Given $h \geq 0$ we have

$$v_k = hv_0 + (1 - h)\delta e_k^T \mathcal{R}\mathbf{v},$$

where $e_k^T \mathcal{R}\mathbf{v} \equiv \sum_{k'=0}^k R_{kk'} v_{k'}$. It follows that we can express the vector of expected payoffs as

$$\mathbf{v} = h\mathbf{v}_0 + (1 - h)\delta\mathcal{R}\mathbf{v} \quad \Rightarrow \quad \mathbf{v} = [\mathcal{I} - (1 - h)\delta\mathcal{R}]^{-1} h\mathbf{v}_0,$$

where we used the fact that matrix $\mathcal{I} - (1 - h)\delta\mathcal{R}$ has full rank, so it is invertible. Hence, the expected payoff if we have today k infected individuals and there is the most extreme lockdown policy (until the disease is completely eradicated) is $v_k \in \mathbf{v}$, which satisfies

$$v_k = e_k^T [\mathcal{I} - (1 - h)\delta\mathcal{R}]^{-1} h\mathbf{v}_0. \quad \square$$

Now we define expected per capita payoffs for general (T, j) policies, i.e., T periods of lockdown followed by j periods of market reopening. Again, let k be the number of individuals infected the period before the lockdown comes into effect. The present discounted value of the average income generated until the policy expires (in period $T + j$) is defined by:

$$v_k(T, j) = \sum_{t=1}^{T+j} \delta^{t-1} (1 - h)^{t-1} h \frac{1 - \delta^{T+j-t+1}}{1 - \delta} \times \bar{y} \\ + (1 - h)^T \delta^{T-1} \sum_{\ell=1}^j \delta^{\ell} (1 - h)^{\ell} \left(1 - \frac{\mu_k(T, \ell)}{N} s \right) \bar{y}.$$

For computational convenience, the calculation does not include incomes past period $T + j$. The first summation calculates the payoff (in present discounted terms) if herd immunity is attained at any point in time $t = 1, \dots, T + j$. This can happen in the lockdown phase, or during the reopening phase, comprising periods $\ell = T + 1, \dots, T + j$. The probability that herd immunity is achieved in period t is $(1 - h)^{t-1} h$. At that point, which we discount by δ^{t-1} , markets reopen. This gives the representative individual a payoff \bar{y} in each subsequent period $t + 1, \dots, T + j$; in period $t + 1$, the present value of this stream is $\frac{1 - \delta^{T+j-t+1}}{1 - \delta} \times \bar{y}$.

The second summation focuses on the reopening phase, i.e., periods $\ell = T + 1, \dots, T + j$, which is why it is discounted by δ^{T-1} . The summation calculates the payoff if herd immunity is *not* achieved. In that case, $\mu_k(T, \ell)/N$ defines the share of infected individuals in the population in a period ℓ . Since a fraction s of these individuals is symptomatic and cannot trade, the upper bound on the income generated in that period is $\left(1 - \frac{\mu_k(T, \ell)}{N} s \right) \bar{y}$.¹⁶ We now proceed by modeling the costs of the disease, current and expected.

¹⁶ It is an upper bound because we are including in the income calculation the income of those who missed their trade meeting because the matching process assigned them to a symptomatic individual. It follows that the lower bound on our welfare calculation is obtained by doubling the loss of income due to the inactivity of symptomatic individuals (i.e., as if no symptomatic was assigned to another symptomatic individual).

Healthcare costs of the epidemic. Our working assumption is that the healthcare cost of the infection depends quadratically on its severity. To measure the severity of the infection we use the share of the population infected in a period. For concreteness, think of the cost generated from the necessary healthcare equipment and resources in place to fight the infection (e.g., setting up and staffing additional medical facilities).

To calculate this expected number, we use Lemma 4. Given that at the end of last period we have $k = 0, 1, \dots, N$ infected individuals, then the expected number of infected individuals at the end of the first period of lockdown is $\hat{\mu}_k(1, 0)$, it is $\hat{\mu}_k(2, 0)$, at the end of the second period of lockdown and it is $\hat{\mu}_k(T, t)$, t periods after the economy has reopened. We divided this number by N to find the expected share of the infected population.

Given k , the per-capita healthcare expenditure is the present discounted value of the cost associated with the policy (T, j) is

$$c_k(T, j) = \sum_{t=1}^T \delta^{t-1} (1-h)^t \left(\frac{\mu_k(t, 0)}{N} \right)^2 + (1-h)^T \delta^{T-1} \sum_{t=1}^j \delta^t (1-h)^t \left(\frac{\mu_k(T, t)}{N} \right)^2.$$

To understand this expression, fix the number k of individuals who are infected before the lockdown policy comes into effect. For computational convenience, the expenditure is only calculated up to period $T+j$, and does not include cost calculations past period $T+j$. In the first phase (T periods) there is no market activity so there is no additional contagion—only recoveries are possible. In each period the disease can be eradicated with probability h , if herd immunity is achieved. In this case, the healthcare costs drop to zero. Hence, in the expression above we only see the terms multiplied by $(1-h)^t$, i.e., the periods $t = 1, \dots, T$ in which herd immunity is not yet attained. In this lockdown phase the expected number of infected individuals is $\mu_k(t, 0)$. This explains the first summation on the RHS of the expression above.

The second summation refers to the reopening phase, i.e., periods $t = 1, \dots, j$ post-lockdown. By the time the lockdown is lifted, there is probability $(1-h)^T$ that herd immunity has not been yet achieved. In that case, $\mu_k(T, t)$ defines the expected number of infected individuals. Since the reopening phase starts T periods after the beginning of the lockdown, we discount the second summation by δ^{T-1} .

Social welfare. Now that we have both the expected cost and the expected income during the $T+j$ periods of the lockdown policy horizon, we offer a measure of average ex-ante welfare:

$$w_k(T, j) := v_k(T, j) - c_k(T, j).$$

It is simply the present value of the per-capita income expected in periods $1, \dots, T+j$ by the average individual, minus the per-capita cost. We use this expression to study the dynamics of social welfare based on different policies (T, j) . The results are reported in the following section. We emphasize that these calculations are meant to offer a numerical illustration, not a carefully calibrated assessment of social welfare. In particular, we emphasize that the welfare measure only considers the trade-off between healthcare costs and lost incomes, and focuses only on the $T+j$ periods that encompass the policy duration. This is done to offer an insight into the tradeoff, while reducing the complexity of the numerical exercise.

5.1. Optimal policy: numerical illustration

We study the welfare impact of lockdowns by means of numerical simulations. For computational convenience, we work




Fig. 8. Welfare gains as a function of lockdown duration. **Notes:** Horizontal axis: duration T of the lockdown phase. Vertical axis: welfare gain relative to a no-lockdown policy. Each point is calculated for a 100 period horizon, with the reopening phase lasting $100 - T$ periods. The circle markers identify the intervention that maximizes welfare. For the medium transmissibility case, imposing no lockdown maximizes the welfare gain for all initial infection rates considered. For the high transmissibility case, imposing some lockdown maximizes the welfare gain with the maximum reached for $T = 1, 3, 6$ as the initial infection rate increases.

with $N = 100$ individuals. All other parameters are as in the earlier experiments.

Fig. 8 reports the results for the cases of medium and high transmissibility, which are the ones that matter in considering the usefulness of lockdowns, give the earlier observations. Three cases are considered, corresponding to three initial rates of infection, 1%, 3% and 10%. Each line illustrates the gain in ex-ante welfare associated with a lockdown lasting $T = 0, 1, 2, \dots, 15$ periods, as opposing to no lockdown. Hence, all curves start from zero. For each T , ex-ante welfare is calculated for the 100 period horizon, so the reopening phase lasts $100 - T$ periods. A lockdown lasting T periods that increases (lowers) welfare generates an observation above (below) zero.

We consider an economy where the infection has not spread out yet and is confined to a small group in the population (initial infection rate 1%). We also consider economies where the infection has already spread out more widely, to 3% and 10% of the population.

The welfare gain reported on the vertical axis is the change in average ex-ante welfare for any given intervention T relative to no intervention, $\frac{w_k(T, j)}{w_k(0, T+j)} - 1$. The length (or, severity) T of the lockdown intervention, from 0 (no intervention) to 15, is reported on the horizontal axis. The circle markers identify the intervention that maximizes welfare, for each initial infection rate scenario. The results of our analysis can be organized into two main observations.

Observation 6. *The welfare impact of lockdowns depends on the diffusion of the epidemic. Immediately imposing a lockdown might lower social welfare.*

The numerical illustration suggests that lockdowns do not necessarily increase welfare. Welfare declines in the figure when the infection is not easily transmitted, which is when the size of lost incomes exceeds the reduction in healthcare costs from the lockdown. Moreover, lockdowns do not necessarily increase social welfare even when the disease is highly transmissible. In

fact, this depends on the duration of the lockdown. Imposing a six-period lockdown is optimal when the infection rate is 10%. The optimal duration declines to 3 and then 1 period as the initial infection rate decreases to 3% and 1%. This suggests a threshold infection rate might exist below which imposing a lockdown is not welfare-improving. In other words, starting a lockdown immediately upon detection of the infection is not necessarily the best strategy even if the diseases are highly transmissible.¹⁷

In this illustration, welfare gains are non-monotone in the initial infection rate. Consider that the maximum welfare gain from lockdowns is about 0.28% when the infection rate is 1%, it climbs to about 0.83% when the infection rate is 3% and then decreases to about 0.01% for an infection rate of 10%. This non-monotonicity is interesting because although we see that longer lockdowns are optimal as the starting infection rate increases, they are not necessarily more socially beneficial. We can make one more observation.

Observation 7. *Prolonged lockdowns can be harmful to social welfare, even for widespread epidemics. Extreme lock-downs are suboptimal.*

Two considerations support this observation. First, Fig. 8 shows that overshooting the optimal duration T of a lockdown can be quite harmful, especially if the disease is highly transmissible but not widespread. In the illustration, welfare gains are rather small and rapidly turn to losses as we move past the optimal duration. For instance, with 1% infection rate, the optimal duration of the lockdown is $T = 1$ periods. This welfare gain turns to a loss as soon as the lockdown duration exceeds three periods. For higher infection rates welfare gains also eventually turn to losses. All curves monotonically decline beyond $T = 15$. Consequently, extreme lockdowns are suboptimal in all cases.

Second, if lockdowns fall short of the optimal lockdown duration this does not dramatically reduce the welfare gains, because of their nonlinearity when the infection is not initially widespread. In this case, a large share of the maximum welfare gain can be attained even if the lockdown falls short of the optimal duration. We also observe that lockdowns shorter than optimal durations can reduce social welfare, when the infection is already widespread. This is because in that case the lockdown does little to rapidly contain the disease and simply reduces current incomes. This creates a welfare loss. However, notice that the welfare reduction from a shorter-than-optimal lockdown is also quite smaller than the welfare loss from a longer-than-optimal lockdown, which again suggests that extreme lockdowns are not optimal. Seen this way, the insight from these numerical illustrations is that it may be best to err on the side of caution, implementing a shorter rather than longer lockdown. On the one hand the risk of overshooting the optimal duration carries a risk of lowering welfare. On the other, falling short of the optimal target does not dissipate too much of potential welfare gains.

An additional aspect suggests a further reason to avoid prolonged lockdowns, or entirely avoid them. The model assumes that the income decline associated with a lockdown is only temporary and as soon as the lockdown is lifted, market activities fully and immediately resume. However, recent experience suggests that this assumption might be empirically unreasonable: some economic activities may be permanently damaged and may not recover when the lockdown is lifted. Below, we enrich the model by introducing the possibility that lockdowns may permanently damage markets to some extent.




Fig. 9. Welfare impact when lockdowns induce permanent damage. **Notes:** Welfare is calculated using expression (5). The parameters are the same as in Fig. 9.

Permanent damage to economic activity. In this section we augment the model, by assuming that a fraction $\alpha(T) \in [0, 1]$ of per-capita income \bar{y} is permanently lost as a consequence of a lockdown. This fraction is an increasing function of the duration of the lockdown, T .

$$w_k(T, j) = v_k(T, j) - c_k(T, j) - \alpha(T) \times \frac{1 - \delta^{T+j}}{1 - \delta} \times \bar{y}. \quad (5)$$

In the numerical illustrations we use $\alpha(T) = (T - 1)/(10T)$. Consider a lockdown lasting 5 periods. This induces a permanent loss of a fraction $\alpha(T) = 0.1$ of per-capita income, which in present-value terms amounts to $\frac{1 - \delta^{100}}{1 - \delta} \times \bar{y}$.

Fig. 9 shows that when lockdowns have permanent adverse consequences on economic activity, their desirability further declines, from a social welfare perspective. This can be ascertained by comparing Fig. 9 to Fig. 8. First, the optimal length of lockdowns decreases even if the epidemic is widespread and highly transmissible. While in Fig. 8 the optimal lockdown lengths increased in the infection rate, in Fig. 9 the optimal length does not exceed one period and, in fact, is zero for 10% rate. Second, the welfare gain now rapidly turns into a welfare loss if the lockdown duration overshoots the optimal target. While in Fig. 8 the welfare gain turned to a loss after five periods lockdown for an initial infection rate of 3%, now welfare declines immediately, as we get into the second period of the lockdown. Overall, this is a further indication supporting Observation 7, i.e., the numerical illustrations suggest care in implementing strict lockdown policies because prolonged lockdowns may end up lowering social welfare, even when the disease is widespread.

6. Discussion

This study contributes to the debate on how to best address the challenges stemming from contagious diseases, such as COVID-19. It offers a mathematical framework – based on the theory of random matching – which makes explicit how economic activity can contribute to the contagion process. The model is used to assess the welfare consequences of non-pharmaceutical interventions that limit economic and social activities—the so-called “lockdowns”. In our model, shutting down all business activity is assumed to be the only way to stop the progress of contagion. Lockdowns can thus reduce the burden on the healthcare

¹⁷ We thank an anonymous Referee for raising this point.

system, but do so by shutting down all income streams as well. The optimal policy intervention must therefore balance these two aspects.

Three insights emerge from the analysis, that seem relevant for policy. First, there is the question of whether imposing drastic stay-at-home mandates at the onset of the disease is the best course of action. For example, Sweden was criticized by the popular press for having kept businesses and schools open. The analysis suggests that there can be gains from imposing some limits to economic activity at the onset of the epidemic because this can be helpful to bring the epidemic under control (Observations 1–4). However, such a policy becomes counterproductive as restrictions on business activity are prolonged, because long-lasting lockdowns are unnecessary to contain the epidemic (Observations 5 and 7).

Second, there is the question of whether all countries should adopt the same response pattern. Our simple model indicates that lockdowns are not a “one-size-fits-all” type of policy. The severity of the intervention should depend on the economic structure in terms of costs imposed on the healthcare system, and the anticipated evolution of the disease based on the characteristics of the susceptible population (e.g., the age structure of a country). All else equal, while stricter interventions might be suitable for countries with more fragile healthcare systems, this policy is not globally optimal, especially in the initial stages of an epidemic if the disease is not highly transmissible (Observation 6). In other words, the model does not lend support to the view that interventions should be necessarily identical across countries. By means of example, supposing that the extreme intervention adopted by China was locally optimal, the model does not imply that such an intervention is optimal for the rest of the world.

In fact, and this is a third insight, extreme and open-ended lockdowns are not generally socially optimal (Observation 7). The numerical illustrations show a nonlinear social welfare response to policy intervention. Tighter restrictions on business activity based on the number of cases does not necessarily constitute an optimal intervention. As policymakers tighten the noose on the economy, they attain progressively smaller healthcare benefits while generating progressively larger income losses. This result clearly depends on the dynamics of healthcare costs, which are assumed quadratic in the share of the infected population, in the numerical experiments presented here.

The numerical experiments reported in this paper should be taken for what they are—illustrative of possible outcomes. They should not be taken as statements having general validity. Indeed, we do not think the model can be practically useful for policy implementation at this stage of its development, as it is still rather rudimentary in more than one aspect. We consider it a first, albeit imperfect, attempt at integrating the contagion process into a richer economic model of the spread of infectious diseases. The question is thus how robust are the insights from this rudimentary model to changes in the model assumptions. Does the proposed model bias the results in favor of or against restrictive interventions and business closures?

On the one hand, the model assumes that business shutdowns, no matter how long, can only induce temporary reductions in income flows. It does not account for possible negative externalities associated with prolonged economic inactivity, such as human capital decline, economic inequality, disruption of business and financial networks, or declines in firm survival rates—phenomena that can depress economic activity for many years to come. This is an especially important consideration for countries whose economic systems are fragile (e.g., consider Greece as compared to Germany). When we enrich the model with the possibility that protracted lockdowns permanently harm economic activity, then even lockdowns of short duration can reduce welfare. Furthermore, in the model business activity is assumed to

spread the infection, and lockdowns prevent further transmission. Yet, there is neither conclusive evidence nor consensus on the empirical reduction in COVID-19 transmission and deaths that can be attributed to lockdowns, or their overall efficacy in reducing case growth (e.g., Allen, 2021; Atkeson et al., 2020; Bendavid et al., 2021; Berry et al., 2021). Additionally, the model assumes that asymptomatic individuals are undetectable infection vectors who have no incentive to avoid market interaction when in fact this problem is greatly mitigated by wide access to quick testing procedures. These assumptions suggest a bias in favor of restrictive lockdowns.

On the other hand, our social welfare calculations consider only a basic tradeoff: that between health costs directly associated with the care needed by infected individuals, and incomes lost due to business closures. There are additional aspects of the COVID-19 disease that are omitted from the model and, hence, bias welfare results in opposite ways. Consider, for instance, the loss of life directly caused by the disease, and the long-term negative health consequences suffered by some COVID-19 patients. Furthermore, we did not account for the beneficial income-smoothing effect of fiscal policy—many countries swiftly addressed the lockdown-induced income decline with significant government transfers financed by an increase in public debt. These omissions bias the social welfare calculations against adopting early and possibly longer lockdowns. This being said, we also ruled out indirect harmful health effects of lockdowns, e.g., increase in mortality from lack of care for other diseases, mental health problems and increased suicides (e.g., see the discussion in Bendavid et al., 2021), and the possibility that outpatient treatments exist that could be effective at reducing death from and COVID-19 hospitalizations (McCullough et al., 2021). Omitting these aspects biases our welfare calculation in favor of lockdowns.

Overall, these considerations suggest that our analysis, albeit rudimentary, is more likely to be biased in favor of finding a beneficial role for lockdowns than a negative one. Future refinements of the model should relax the assumptions discussed above to enhance the empirical applicability of the model. Other features missing from the current layout should also be included. For instance, our current formulation of the severity of the intervention is only governed by the duration parameter T . One could consider an additional dimension that accounts for the proportion of business activity affected by stay-at-home mandates. Another extension is to consider contagious processes that are affected by the population density, something that can be implemented by assuming the contagion parameter p depends on the population size N . Here, we have chosen to keep the framework nimble in order to lay out as clearly as possible an explicit process of contagion – our pairwise random matching process, that is – and show how it can be integrated into economic models of epidemics.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M., 2020. Optimal targeted lockdowns in a multi-group SIR model. NBER working paper no. 27102.
- Allen, D.W., 2021. Covid lockdown cost/benefits: A critical assessment of the literature. Unpublished manuscript, Simon Fraser University.
- Alvarez, F., Argente, D., Lippi, F., 2020. A simple planning problem for COVID-19 lockdown. NBER working paper no. 26981.

- Atkeson, A., Kopecky, K., Zha, T., 2020. Four stylized facts about covid-19. NBER working paper 27719.
- Avery, C., Bossert, W., Clark, A., Ellison, G., Ellison, S.F., 2020. An economist's guide to epidemiology models of infectious disease. *J. Econ. Perspect.* 34 (4), 79–104.
- Bendavid, E., Oh, C., Bhattacharya, J., Ioannidis, J., 2021. Assessing mandatory stay-at-home and business closure effects on the spread of Covid-19. *Eur. J. Clin. Invest.* <http://dx.doi.org/10.1111/ECL.13484>.
- Berry, C.R., Fowler, A., Glazer, T., Handel-Meyer, S., MacMillen, A., 2021. Evaluating the effects of shelter-in-place policies during the COVID-19 pandemic. *Proc. Natl. Acad. Sci. USA* 118 (15), <http://dx.doi.org/10.1073/pnas.2019706118>.
- Busenberg, S., van den Driessche, P., 1990. Analysis of a disease transmission model in a population with varying size. *J. Math. Biol.* 28, 257–270.
- Camera, G., Giofr , A., 2014. A tractable analysis of contagious equilibria. *J. Math. Econom.* 50, 290–300.
- Centers for Disease Control and Prevention, 2020. SARS-CoV-2 pandemic planning scenarios. Retrieved on Jan 4, 2021 at <https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html>.
- Diamond, Peter, 1982. Aggregate demand management in search equilibrium. *J. Polit. Econ.* 90, 881–894.
- Eichenbaum, M., Rebelo, S., Trabandt, M., 2020. The macroeconomics of epidemics. NBER working paper no. 26882.
- Goenka, A., Liu, L., Nguyen, M.-H., 2014. Infectious diseases and economic growth. *J. Math. Econom.* 50, 34–53.
- Ioannidis, John P.A., 2021a. Reconciling estimates of global spread and infection fatality rates of COVID-19: An overview of systematic evaluations. *Eur. J. Clin. Invest.* 51, 1–13. <http://dx.doi.org/10.1111/eci.13554>.
- Ioannidis, John P.A., 2021b. The infection fatality rate of COVID-19 inferred from seroprevalence data. *Bull. World Health Org.* 99, 19–33F. <http://dx.doi.org/10.2471/BLT.20.265892>.
- Kory, P., Meduri, G.U., Varon, J., Iglesias, J., Marik, P.E., 2021. Review of the emerging evidence demonstrating the efficacy of ivermectin in the prophylaxis and treatment of COVID-19. *Am. J. Ther.* 28 (3), e299–e318. <http://dx.doi.org/10.1097/MJT.0000000000001377>, (2021).
- McCullough, P.A., Kelly, R.J., Ruocco, G., Lerma, E., Tumlin, J., Wheelan, K.R., Katz, N., Lepor, N.E., Vijay, K., Carter, H., Singh, B., McCullough, S.P., Bhambhani, B., Palazzuoli, A., De Ferrari, G.M., Milligan, G.P., Safder, T., Tecson, K.M., Wang, D.D., McKinnon, J.E., O'Neill, W.W., Zervos, M., Risch, H.A., 2021. Pathophysiological basis and rationale for early outpatient treatment of SARS-CoV-2 (COVID-19) infection. *Amer. J. Med.* 134 (1), 16–22. <http://dx.doi.org/10.1016/j.amjmed.2020.07.003>.
- Mortensen, D., Pissarides, C., 1994. Job creation and job destruction in the theory of unemployment. *Rev. Econom. Stud.* 61 (3), 397–415.
- The White House, 2021. Remarks by president Biden on the COVID-19 response and the vaccination program. Briefing Room, Speeches and remarks, May 04. <https://www.whitehouse.gov/briefing-room/speeches-remarks/2021/05/04/remarks-by-president-biden-on-the-covid-19-response-and-the-vaccination-program/>. (Accessed 15 May 2021).
- Wall Street Journal, 2021. Germany and Israel tighten lockdowns to buy time for vaccinations. By James Hookway, Jan 6, 2021. <https://www.wsj.com/livecoverage/covid-2021-01-06/card/t2ez497yO5x5MauXUf65>. (Accessed 13 May 2021).
- World Health Organization, 2020a. Coronavirus disease (COVID-19): Similarities and differences with influenza. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/question-and-answers-hub/q-a-detail/coronavirus-disease-covid-19-similarities-and-differences-with-influenza>. (Accessed 24 January 2021).
- World Health Organization, 2020b. Coronavirus disease (COVID-19): Herd immunity, lockdowns and COVID-19. <https://www.who.int/news-room/q-a-detail/herd-immunity-lockdowns-and-covid-19>. (Accessed 13 May 2021).